DOCUMENT RESUME

ED 235 476 CS 007 333

AUTHOR Ortony, Andrew; Radin, Dean I.
TITLE SAPIENS: Spreading Activation Processor for Information Encoded in Network Structures. Technical Report No. 296.
INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
SPONS AGENCY National Academy of Education, Washington, D.C.; National Inst. of Education (ED), Washington, DC.
PUB DATE Oct 83
CONTRACT 400-76-0116
NOTE 39p.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Cognitive Processes; *Computational Linguistics; *Computer Programs; Computers; Computer Science; *Information Retrieval; Language Processing; Language Research; Reading Research
IDENTIFIERS Relevance (Evaluation)

ABSTRACT

The product of researchers' efforts to develop a computer processor which distinguishes between relevant and irrelevant information in the database, Spreading Activation Processor for Information Encoded in Network Structures (SAPIENS) exhibits (1) context sensitivity, (2) efficiency, (3) decreasing activation over time, (4) summation of activation, and (5) an activation threshold for each node that determines whether it will transmit activation to other nodes. Used for natural language processing, when given several input words, SAPIENS can quickly identify from its 16,000 word database a set of 10 to 20 relevant words without extensive searching. SAPIENS could be employed in other artificial intelligence domains as well. (MM)

**********
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
**********

ED235476

CENTER FOR THE STUDY OF READING

Technical Report No. 296

# SAPIENS: SPREADING ACTIVATION PROCESSOR FOR INFORMATION ENCODED IN NETWORK STRUCTURES

Andrew Ortony
University of Illinois at Urbana-Champaign

Dean I. Radin
Bell Telephone Laboratories, Columbus, Ohio

October 1983

University of Illinois
at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02238

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

The research reported herein was supported in part by the National Institute of Education under Contract No. US-NIE-C-400-76-0116, and in part by a Spencer Fellowship awarded by the National Academy of Education to the first author. We wish to thank George Kiss for supplying a magnetic tape version of the Associative Thesaurus, the Coordinated Science Laboratory and Computing Services Office of the University of Illinois for granting computer time on the DEC-10, David Waltz for the use of a disk pack which housed a modified version of the Associative Thesaurus, and Harry Blanchard and Glenn Kleiman for their helpful comments on earlier versions of this manuscript.

CS 007333

EDITORIAL BOARD

William Nagy
Editor

Harry Blanchard
Nancy Bryant
Pat Chrosniak
Avon Crismore
Linda Fielding
Dan Foertsch
Meg Gallagher
Beth Gudbrandsen
Patricia Herman
Asghar Iran-Nejad
Margi Laff
Margie Leys
Theresa Rogers
Behrooz Tavakoli
Terry Turner
Paul Wilson

Abstract

This report describes a computer implementation of a spreading activation process in semantic memory and discusses its performance on some tasks often employed in psychological studies of human language processing. An associative thesaurus containing over 16,000 words and all free-associative strengths between them was used as the data base, thus making SAPIENS confront the information and computation problems inherent in large data base manipulation.

# SAPIENS: Spreading Activation Processor for Information Encoded in Network Structures

One of the hallmarks of intelligence is the ability to efficiently identify and utilize information relevant to the solution of a problem while largely discounting irrelevant information. In many cognitive tasks (such as language comprehension or perception) this means that relevant information must be accessed from memory more or less immediately, thus precluding any kind of exhaustive, or near exhaustive search. Consequently, the relevance problem--the problem of how to identify a potentially relevant subset of the totality of stored information--is an important theoretical question in psychology and an important practical question for AI (Artificial Intelligence). The work described in this report takes an AI perspective on the problem, using the context of natural language processing as its basis, although the principles upon which it is based are quite general.

Research in AI has devised various domain-specific mechanisms for dealing with the relevance problem. Robinson's (1965, 1968) resolution principle is an early example of an approach that proved fruitful in the field of automatic theorem proving. In the domain of scene analysis Waltz (1975) solved the problem by taking advantage of the huge reduction in data storage resulting from distinguishing the physically possible pairs from the logically possible pairs of line junctions in a two dimensional representation of a scene. In the area of problem solving proper, a general guiding principle has been the careful choice of knowledge representation. It quickly became apparent that the choice could have a dramatic influence on the ease of problem solution, the old problem

of the mutilated chess board being a simple but convincing example (see, for example, Raphael, 1976).

Two developments in AI and psychology are especially pertinent to the relevance problem. The first is the emergence from earlier but vaguer accounts (e.g., Bartlett, 1932; Piaget, 1952) of increasingly detailed proposals about the nature of generalized knowledge representations, variously called frames (e.g. Minsky, 1975), scripts (e.g. Schank & Abelson, 1977), and schemas (e.g. Rumelhart & Ortony, 1977). The second, related, development, primarily from AI, is the recognition that the distinction between programs and data need be much less sharp than was generally supposed. The notion is that much information that appears on the surface to be declarative in nature can be, and often is more advantageously represented as procedural. This observation constitutes an important underlying principle of Planner-like languages (c.f. Hewitt, 1972), and is also evidenced in the system of Norman and Rumelhart (1975). A consequence of the devaluation of the program/data distinction is that generalized knowledge structures, here to be called "schemas," are partly procedural and partly declarative.

Because the utilization of a schema (in an AI system or in human cognition) results in a great deal of potentially relevant information becoming available automatically, it offers a powerful way of dealing with part of the relevance problem, but it does only deal with part of it. The missing component is the schema selection mechanism which is responsible for bringing the appropriate candidate schemas into play in the first place. In this report we describe a computer implementation of a process for doing this--a proc roposed earlier in Ortony (1978).

Efficient access to all and only the relevant knowledge structures depends primarily not so much on the internal structure of individual knowledge representations (which is the problem of chief concern to those working on scripts, schemas, etc.) as on the overall organization and interrelations of such representations in the system. In other words, it depends on the overall structure of memory, rather than upon the the structure of the things represented in memory. Models of the macrostructure of knowledge organization have tended to rely upon associative networks. However, associative networks (whether implemented or not) have generally been viewed merely as spaces in which to conduct a specific kind of search operation known as the intersection search (Quillian, 1968; Collins & Quillian, 1969; Collins & Loftus, 1975). The general mechanism employed is that of spreading activation, and the mechanism is considered to have succeeded when it discovers an intersecting node that can be reached from the different source nodes. This limited use, however, fails to capitalize on the power both of network representations themselves, and of the spreading activation mechanism. The principal purpose of SAPIENS was to harness the potential power of the spreading activation mechanism and the semantic network representation to simulate schema selection. Furthermore, this was undertaken in the context of a data base of sufficient size that the principles of the system's design could be generally applicable rather than relegated to the category of ungeneralizable "toy" problems. Given this goal (as opposed to that of schema utilization), it was possible to ignore the structure of the nodes in the net. The nodes are simply English words, although we make the assumption that as such, they can, in principle, be treated as the names of schemas.

Our assumption that a network of words can be regarded as a netweork of schema names is an important simplifying assumption which warrants some

elaboration. In a complete representation, we assume that there would have to be at least three different levels--a lexical level, a conceptual level, and an episodic level. The lexical level represents connections between words without distinguishing between distinctly different meanings that a word might have. At the lexical level a word like bank could have connections to other words, some of which (e.g., money) have to do with the "financial institution" sense, some (e.g., weeds) with the "side of a river" sense, and some with other senses of the word such as those related to basketball, airplanes, etc. Thus, the lexical level represents associations between words, not between concepts. Conceptual connections are represented at the next, conceptual, level. At this level, we suppose that there are separate schemas for the distinct meanings of a word like bank. Furthermore, these schemas are ordinarily not directly connected to one another. However, they are connected through the lexical level in the sense that they are all directly connected to the word or words that constitute their labels or names. Finally, we assume that representations involving some of the more noteworthy specific experiences centered around particular schemas are represented at the episodic level. At this level, individual representations are again directly associated with particular schemas, perhaps indexed in terms of notable deviations from the canonical representations (see Schank, 1982). In the present report, we investigate the degree to which schema selection can be facilitated through processes that are restricted to the lexical level. There are both practical and theoretical reasons why this is a worthwhile enterprise. The practical reason is that it is much easier to construct a data base of lexical associations than it is to construct a comparably sized data base of conceptual and episodic structures. The theoretical reason is that the schema selection process is of necessity a pre-semantic process. That is, it is a process whose goal is to permit a determination of the meaning of some input.

Consequently, the process cannot presuppose that a semantic representation of the input has already been achieved. This means that the schema selection process has to operate at a relatively impoverished semantic level. Of the three levels that we have proposed, only the lexical level is devoid of semantic content.

A spreading activation mechanism ought to have at least the following five characteristics: (a) context sensitivity, which permits the production of different patterns of activation for the same input string under different context conditions; (b) efficiency, permitting a mechanism to operate in a space containing perhaps tens of thousands of nodes; (c) decreasing activation over time, so as to prevent every input from activating the entire network forever; (d) summation of activation from different sources, so as to permit differential activation levels on equally distant nodes; and (e) an activation threshold for each node which determines whether it will transmit activation to other nodes.

Based on these principles, a processor for operating on a network was developed. In it, spreading activation is used as a mechanism not just for finding intersections, but for identifying constellations of candidate nodes for employment in the process at hand. In other words, the mechanism is used to restrict the set of potentially relevant nodes. The ability to select a small set of of potentially relevant nodes for possible use in subsequent processing is, as we have already suggested, an important component in an intelligent system. While we acknowledge that it is not sufficient to endow a system with genuine wisdom, we were unable to resist the name SAPIENS--Spreading Activation Processor for Information Encoded in Network Structures. The program was written in MACLISP and implemented on a DEC-10 computer. The data base was an associative thesaurus consisting of over 16,000 words and all free-associative strengths between them (Kiss, 1968). Since the data base was empirically

generated (by soliciting associative responses from students), and since it is quite large, the network can be regarded as a reasonable facsimile of (part of) some arbitrary individual's lexical map.

It was considered important to use a large, empirically realistic data base for two reasons. First, in order to be able to test the effectiveness of the simulated process, it was necessary to have a data base with a great deal of semantic diversity and with sufficiently rich interconnections to avoid trivializing the problem. Second, since a large data base was used, the so-called "combinatorial explosion" problem had to be addressed. Most computer simulated semantic network models contain less than a thousand nodes (for example, Quillian, 1968, encoded about 850 nodes). The present system works on a data base an order of magnitude larger than any other semantic network system we are aware of. Clearly the time requirements resulting from the massively increased number of possible paths in a network of 16,000 nodes put much greater demands on the processor. Finally, we felt that only with a relatively large and semantically diverse data base would it be possible to explore the potential of the system for dealing with a broad range of tasks.

The main result of SAPIENS is that, given several input words, it quickly identifies from the entire 16,000 word network a restricted set of 10 to 20 relevant words without extensive searching. These words can be thought of as the names of the best candidate schemas for subsequent top-down processing by other mechanisms. In this respect, SAPIENS is analogous to the filtering process in Waltz's (1975) program that generates semantic descriptions of scenes with shadows. This filtering process takes a scene and finds a small subset of most-likely line segments (out of thousands of possibilities) for further processing by a semantic description mechanism. SAPIENS takes an input string and finds a

small subset of relevant concepts potentially usable for further processing by, for example, inference or problem solving mechanisms.

Four tasks were used to examine the performance of SAPIENS, While some of these tasks were suggested by published experimental work, it is important to emphasize that they are not intended as simulations of experiments. Rather, they should be viewed as illustrations of the kind of problems that SAPIENS can handle and of the way in which it handles them. Thus, we do not view the tasks as merely being relevant to the question of whether a spreading activation model can, in principle, provide simple solutions to the schema selection problem and to such issues as lexical disambiguation. Proposals that some mechanism or other can in principle solve some set of problems are not very compelling. We view performance on the tasks as demonstrating that a spreading activation model employing a realistically large number of nodes does solve these problems. Furthermore, we consider it important that we have a working program to do this, rather than a theoretical proposal--the enterprise, therefore, is an AI enterprise.

The first task shows how context can be used to disambiguate ambiguous words. The second task seeks to show how standard typicality effects found in various laboratory tasks can be accommodated by SAPIENS. Third, SAPIEN's performance in a mock cued recall "experiment" is examined. Finally, we describe a simple examination of the effects of manipulating word order of an input string.

It should be emphasized that we do not claim that the mechanism we propose is sufficient to realize complete solutions to all such problems. Clearly schema utilization is equally important. However, we do claim that, properly conceived, a spreading activation mechanism may well be a fundamentally

important ingredient insofar as it represents a viable solution to the schema selection problem. We consider it significant that so much can be achieved with so little, considering that the only data that SAPIENS operates with are associative strengths between words.

## Design Features

Before discussing the design characteristics of the program it will be helpful to define and clarify our use of a number of key terms. First, in the discussion that follows we do not distinguish between nodes, words, and schemas or concepts. This is not because there is no difference, but because in the present context the differences do not matter. The network contains nodes which are representations of English words. However, our theoretical orientation views words as labels or names for schemas or concepts. We can consistently maintain this position even though we have made no attempt to implement the underlying schematic structure that such nodes would have to have if they were to be employed in a full-fledged natural language processor.

The network itself, then, consists of some 16,000 English words and their links which represent the empirically based associative strengths between them. Associative strength is an integer measure (1-100) of how many times a response (R) was given to a specific stimulus (S) by 100 people (Kiss, 1968). The proportion of the associative strength "sent" from S to R is called activation. The activation level of a word is the total amount of activation currently being received at that word. It is the sum of all activations coming into that word from any other stimulus words that are "transmitting" or "sending" activation. The associative thesaurus contains both forward (S-R) and reverse (R-S) data. The set of nodes that can be reached from (i.e. as responses to) a particular word is called the forward environment.

The words that are used as input strings to SAPIENS are called seeds. Each seed is weighted so that the original associative strengths can effectively be altered to simulate context effects and decreasing activation over time. An expansion consists of accepting each input seed in turn as a stimulus and creating a single list of all responses to these stimuli. After the expansion is complete, intersections are found within the expanded list, and these intersections, along with the activation levels associated with them, comprise the relevant forward environment or RFE. The expansion of all the input seeds and the identification of intersections (and therefore the RFE) occur within one time-slice or spread. This breadth-first expansion in one time-slice simulates a parallel computation process, and thus spreading activation in SAPIENS can be regarded as a parallel process which simultaneously spreads activation out from several input seeds.

The operation of SAPIENS is analogous to growing a crystal in a saturated solution. Input seeds used to start spreading activation are similar to chemical seeds used to start a crystal-growing process, and the densely interconnected associative network is similar to a densely saturated chemical solution. The seeds constitute a core around which layer upon layer of molecules (or nodes) cluster. This clustering produces an onion-like series of shells around the seed core. The first layer is formed of molecules that are most strongly attracted to the seeds. After several layers have formed, there is less of a tendency for the crystal to continue growing. At some further point, an attraction threshold fails to be reached and the growing process stops. In a similar fashion, the first SAPIENS spread creates a cluster of strongly connected nodes around the input seed core. As each new layer is formed, the size of the cluster increases, and more nodes are pulled into the RFE. Each new layer, however, is formed of successively more remotely related nodes.

The first spread creates a tight cluster around the seed core. The second spread wraps another layer about the first, but is associatively less strongly related to the core. Normally, SAPIENS stops after the second spread since successive clusters are mildly related at best, but more importantly, the purpose for the spreading has been fulfilled in one or two spreads, namely, to quickly find a small subset of relevant concepts for further processing. Not all of the words found represent concepts that are likely to be useful for further processing, but the ones that are useful are found without searching the entire network.

On average, each node in the network spreads out to about 20 associates. In a goal-searching algorithm, thousands upon thousands of nodes would eventually be reached, especially since the network is so densely interconnected. It is this same density that enables the clustering approach to prevent untold thousands of nodes from being activated. At each new spread the original seeds plus the intersections are re-input as new seeds. The result of this re-input is to restrict the kinds of intersections that will result. As more and more intersections are input as seeds, it becomes more likely that most of the activation will circle in on itself, that is, an area of very high activation will begin to form as more intersections are added to the inputs. The inputs are recursive, each input containing part of the previous one, and the effect is to quickly excise from the network the relevance structure that exists around the original inputs.

Design Example

We shall now describe a fairly detailed example of the operation of SAPIENS using, for simplicity's sake, fictitious but realistically representative data. Suppose the input seeds are the two words, apple and fruit, both with input weights of 1. First, apple and fruit are expanded to create a list containing all responses and associative strengths from the seeds. It is important to notice that the total activation output from each seed equals 100, as shown in Figure 1. The value 51, for example, from fruit to banana, refers to the fact that the proportion of subjects producing banana as a response to the stimulus fruit was 0.51. In other words, they are empirically-based transition probabilities.

---

Insert Figures 1 and 2 about here

---

The next step is to find intersections in the expanded space and to create a new list of the intersections and summed strengths as illustrated in Figure 2. Once the intersections are found and the strengths summed, a new list is created as input for the next time-slice. The new input list contains the original seeds with input weights 1, together with the intersection words at weights which reflect the proportion of total activation on each word (see Figure 3). The purpose of using proportional weights on the intersected nodes is to prevent an activation explosion on the next time-slice. The weights on a seed are used to multiply the associative strengths of each response to that seed. For example, since the weight on the node red is 0.3, the activation sent to each response from red will be multiplied by 0.3. This means that the total activation sent from red will be 0.3 x 100, or 30, because each node originally outputs a total of 100. The original seeds are always re-input at weight 1 and

successive intersection nodes are always input at proportional weights which sum to 1 to simulate the attention being placed on the input seeds (high activation) and the process of spreading activation (low activation) operating on intersected nodes.

---

Insert Figures 3 and 4 about here

---

Since the total input weight of all intersected nodes is constrained to sum to 1 (so as to prevent an activation explosion), the activation on intersections decreases rapidly. As the new seed list increases in length on each successive spread, the proportional activation on each word in the relevant forward environment (RFE) must decrease. It is, of course, possible, and desirable, that the proportional activations on relevant RFE words will change as the spread continues because some nodes will receive additional activation from newly accessed nodes.

Continuing the example, with the creation of the new input list, the spread cycle (and one time-slice) is complete. At the end of each cycle the result is a new RFE. Notice, as can be seen in Figure 4, that the input seeds apple and fruit do not have strength sums. Normally all nodes that appear in the RFE must have received activation from at least two sources. The exceptions are the input seeds, which are always included in the RFE whether they received activation or not. Each word in the RFE has an activation strength associated with it which is the total amount of activation received by that word in the time-slice that just ended. The activation strength sum is the sum of all activation levels of nodes currently in the RFE. In the present example, the activation strength sum is 74.

The maximum strength sum attainable after any spread is 100xN + 100, where N equals the number of input seeds. This is because each input seed always outputs a strength of 100 (weight = 1), and the sum of strengths from all intersected nodes is forced to equal 100 (sum of weights = 1). In our example, if a total activation of 300 is injected into the network at the beginning of the second spread, then 300 is the maximum attainable strength sum. This can occur only if all of the responses (from the input seeds and the intersections) happen to be intersections too, for these particular responses are the only nodes receiving activation in the network.

A maximum attainable strength sum enables us to make meaningful comparisons among different strength sums. For example, if the maximum strength is 300, and one pair of seeds produces a sum of 60 after 2 spreads, and another pair produces a sum of 30 after 2 spreads, we can say that the first pair represents 20 percent of the available activation and the second pair represents 10 percent. Thus, the first pair of seeds is better "integrated" (in the sense of "inter-related") than the second pair by 10 percent of the available activation.

Input seeds that are strongly associated with one another will produce an RFE containing more intersections and hence a higher activation strength sum. Weakly associated seeds may produce an RFE with only a few intersections, and a correspondingly smaller activation strength sum. "Association" in this sense is thus better thought of as a measure of "inter-relatedness" between words.

## Program Performance

### Lexical Disambiguation

The first task to be described is concerned with the problem of the selection of appropriate word meanings in the face of more than one alternative. The most typical example of this problem is the need to disambiguate homonyms on the basis of contextual information. As J. Anderson (1976) puts it:

> One of the problems with all current [computer-based language comprehension] models is that they run into difficulties when they must deal with multiple word senses or multiple syntactic possibilities ... The spreading activation model provides the potential for associative context to prime a word's meaning. These parallel, strength mechanisms ... are hard to simulate on a computer, but this difficulty is irrelevant to the question of their psychological validity. (p.448)

The degree to which SAPIENS is able to contribute to the solution of disambiguation problem was investigated by injecting a context-setting word and the target (ambiguous) word into the network as input seeds. The intersections found in one or two spreads were usually sufficient to produce clusters of words that were clearly associated with the contextually appropriate meaning of the disambiguous word. For example, consider the words mint, bank, bar, fruit, and game. The first three are ambiguous in the usual sense: they each have more than one distinct meaning. In particular, we were interested in distinguishing between the sense of mint as candy, and as a place for manufacturing coins. Similarly, we were interested in distinguishing the sense of bank as a financial institution from the sense in which rivers have banks. And we wanted to distinguish the sense of bar as a place to have a drink, from the sense in which a rod is a bar. The remaining cases, fruit and game, are not ambiguous words in

the usual sense, rather they are vague, being consistent with a wide range of very different kinds of referents.

If SAPIENS is appropriately sensitive to context it ought to provide evidence of greater availability of words related to the appropriate meaning of an ambiguous word. In the case of the vague terms, context ought to impose constraints on the relative accessibility of different instances of the concept. Thus, for example, if red is part of the context for fruit, apple ought to be more available than banana, or lemon, whereas, if yellow were part of the context, one might expect the reverse to be true. The results of the simulation are presented in Table 1.

---

Insert Table 1 about here

---

Recall that the activation strength sum can be taken as a measure of the degree of interconnectedness or integration of the word clusters associated with an input. Thus, for example,the first two rows in Table 1 suggest that the pecuniary sense of mint is better integrated than the gastronomic sense. Another interesting property of SAPIENS is that the original seeds only appear in the RFE if they themselves receive activation from their associates. So, river appears high in the RFE (see Row 3 of Table 1) because it received activation from water, stream, and/or other activated elements in the RFE.

A general observation that can be made about these results is that the most highly ranked nodes in the RFE are, as a rule, semantically highly related. In cases where the RFE is large (as indicated by a high strength sum), for example bank in the context of money, the least highly ranked concepts are very weakly related, often representing syntagmatically, rather than paradigmatically related concepts. For example, make appears very low on the list for bank.

Presumably it is only there at all because one puts money that one makes into the bank. One way to interpret the varying numbers of nodes in the RFE is to suppose that the more nodes there are in it, the better integrated the RFE is, and the more knowledge there is associated with the particular use of the term (compared only, of course, to alternative uses of the same term). In other words, uses giving rise to larger RFEs might be regarded as the higher frequency, or more typical ones. However, caution should be exercised in this respect: it is not being claimed that, for example, the most probable use of the word mint, is the pecuniary sense, but only that that sense is more probable than the particular alternatives that were investigated in the context of the present data base. It is perfectly possible that mint as herb, would be the preferred sense in some other data base. Furthermore, a realistic test of this would need to control for the frequency of the context-setting word, which was not done in the present case.

If the program is viewed as a model of schema selection (presumably for later use in top down processing), it can be concluded that SAPIENS is reasonably successful. The clustering process was supposed to quickly find a small relevant subset of the nodes without resorting to extensive searching. This relevant subset, the RFE, ought to have, and did include, nodes that were related to the input seeds taken together. The program, it can therefore be concluded, is able to restrict the candidate nodes for subsequent processing so that those candidates are likely to be relevant to the appropriate meaning of a word as determined by the context.

Note that SAPIENS does not know which seed is the context-setting word and which is the target word. If both nodes are somewhat ambiguous, a "complementary disambiguation" can occur. Consider for example, drink and bar. Both words are compatible with several different kinds of referents. The bartender made a

drink and The mother made a drink suggest different kinds of drinks while The bar was crowded and The bar was rusted imply different kinds of bars. Since SAPIENS finds items relevant to all of the input words, the RFE for drink followed by bar will reflect both a bar sense of drink and a drink sense of bar. Similarly, finding a relation between red and fruit, and fruit and red entails a comparable process, except that the seeds are (technically) unambiguous.

Typicality and Verification Tasks

It is now a well-established finding in Psychology that the speed with which true sentences of the form "An x is a y" can be verified depends on how typical the example, x, is of the category, y, (see, e.g., Rosch, 1973; Smith, Shoben, & Rips, 1974). If apple is a more typical fruit than strawberry, then the sentence An apple is a fruit will take less time to verify than the sentence A strawberry is a fruit. This fact is, of course, consistent with the view that more typical exemplars are more available than less typical ones. The second set of simulations was conducted to see whether or not SAPIENS would demonstrate such differential availability.

Injecting apple followed by fruit into the network results in an RFE containing the words pie, pear, orange juice, tree, juicy, tart, banana, green, sauce, and summer after just one spread. These words are listed in order of activation level; the proportion of available activation was 0.32 (sum strength for this RFE was 97). By comparison, if the seed nodes are strawberry and fruit, the RFE comprises cream, juice, jam, raspberry, pie, tart, summer, and food after one spread, with an activation proportion of 0.21 (sum strength equals 64). Thus, as expected, the more typical exemplar, apple, produces an RFE with greater overall strength than the less typical one. Similarly, robin followed by bird gives bird, song, sparrow, swallow, thrush, and starling after

two spreads, with an activation proportion of 0.38 (sum strength of 114), while penguin followed by bird gives bird, fly, black, feathers, feather, flight, sky, and wing after two spreads, with a smaller activation proportion of 0.29 (sum strength of 87) than in the robin case.

At this point it is worth making a couple of observations about the contents of RFEs. First, it should be remembered that the data that form the basis of the network were collected several years ago and represent some amalgam of a large number of undergraduate students at Edinburgh, Scotland. Assuming the general tendency of the dialect to be closer to that of British English than to American English, it should be recognized that robins (being relatively uncommon in Britain) are not, in fact, the best exemplars of birds--sparrows, thrushes, and starlings are all better. Notice that these other good instances in fact find their way into the RFE for robin and bird. There are a number of other peculiarities of this kind. For example, it is the experience of most English people that apples are typically just as green as they are red (perhaps even more so) because there is a subcategory of apples known as cooking apples which are always green (and sour) and are used in pies and tarts. Another peculiarity by U.S. standards might be the occurrence of pint in the RFE for drink bar. In Britain one of the most typical things to drink in a pub (especially at the bar) is a pint (of beer).

Notice also, that the RFE often contains more than one cognate as a word (e.g. feather and feathers). Sometimes this is just a plural form, sometimes a verb form, and so on. Ideally, these would have been removed, but their presence causes no real problems. Furthermore, it is interesting to note that the second highest RFE word for the pair penguin/bird was fly. It is possible that words like feathers and wing contributed enough activation to fly for it to become highly activated. On the other hand, perhaps fly refers to the fact that

penguins are unable to fly. As it is currently set up, SAPIENS cannot handle negative properties, but again, for present purposes this is not really important. Finally, it should be mentioned that psychological experiments on sentence verification routinely reveal not only that subjects are faster at verifying good examples of category members than poor ones, but also that they are very fast at rejecting non-examples like A pencil is a bird. The RFEs resulting from the corresponding inputs can be interpreted in a manner consistent with this finding, for, in the case where the term refers to a non-instance of the category, the RFE tends to be empty so that the sum strength is equal to or close to zero.

Cued Recall

In an experiment reported by Anderson and Ortony (1975), subjects studied a number of simple sentences such as (1) and (2).

(1) Nurses are often beautiful.

(2) Nurses have to be licensed.

Later, subjects were given either a "close" or a "remote" recall cue. In the present example actress would be the close cue for sentence (1) and the remote cue for sentence (2), while doctor would be the close cue for sentence (2) and the remote cue for sentence (1). The experiment showed that what was semantically close or remote depended on the entire sentence rather than on part of it since the cues were not sufficient to permit the recall of control sentences which included key words like beautiful and licensed. For example,

the cues were not differentially effective, indeed, they were not effective at all for control sentences like (1a) and (2a).

(1a) Landscapes are often beautiful.

(2a) Taverns have to be licensed.

Ortony (1978) has suggested that a model along the lines of the present one can account for these results. Different RFEs are created for the different sentences: close recall cues integrate strongly with their respective RFEs, and distant recall cues integrate weakly with them. These integration strengths reflect the probability of a cue eliciting recall of a sentence.

The task in the Anderson and Ortony (1975) experiment was simulated by first creating RFEs from input seeds corresponding to the substantive words in the to-be-learned sentences. Then the expanded cues were intersected with the input seed RFEs to find the amount of integration between them. The sequence of operations used by SAPIENS is illustrated in Figure 5.

---

Insert Figure 5 about here

---

The process begins when seeds A and B are injected into the network (Figure 5a). Intersections are then found between expanded A and expanded B. This area is called RFE1 (Figure 5b). The seeds and RFE1 are re-input to the network (Figure 5c) as explained in the basic design section. This gives rise to a second RFE (RFE2) resulting from the intersections of seeds A and B, with RFE1 (Figure 5d). Any word that receives activation from at least two sources is part of this (and subsequent) RFEs because the notion of a threshold requires that a certain amount of activation is received by a node before it begins to transmit activation to other nodes. In SAPIENS, a node receiving activation from

just one source will not exceed its threshold because insufficient activation will be received (Figure 5d). The intersection between the RFE2 and the expanded cue is an example of a seed-cue intersection (Figure 5e). The activation strength sum of the seed-cue intersection is directly related to the probability of recalling the input given a cue because the activation strength sum is a measure of the overlap between the input seed RFE and the cue RFE. It is reasonable to assume that if a cue RFE completely overlaps (and restimulates) the input seed RFE that the probability of recalling the inputs will be very high. On the other hand, if a cue RFE does not overlap any portion of the input seed RFE, it is unlikely that the input will be recalled because no close associates of the inputs were reactivated. Therefore, if one activation strength sum is greater than another, it seems reasonable to assume that the former sum will reflect a greater probability of input sentence recall than the latter sum. The results are summarized in Table 2.

---

Insert Table 2 about here

---

In almost all cases the trends are in the direction found by Anderson and Ortony (1975). Not all the words from the original experiment were available in the network, consequently some minor changes were needed. None of these wording changes in any way affected the validity of the test. For example, the use of the cue sexy, for the sentences about nurses was necessitated by the fact that the original cue used in Anderson and Ortony (1975), actress, was not in the network. As an aside, it should perhaps be mentioned that this forced us to revert (for the purposes of science, only) to the standards of sexism that were pervasive at that time the data base was assembled some 20 years ago.

It is important to emphasize that no claims are being made here about specific comprehension and memory processes. All that is being supposed is that whatever the comprehension process is, it does involve the establishment of an RFE to limit the schemas that are brought into focus for that process. It is also assumed that that RFE could be used as part of the representation of the remembered sentence (or reproduced from it), a part that apparently can be gainfully employed in the retrieval process.

Word Order Effects

Compare sentences (3a) and (3b).

(3a) The boy smashed the bus.

(3b) The bus smashed the boy.

In sentence (3a) images of vandalism and boy-related items are likely to come to mind, whereas in sentence (3b), accidents and bus-related items come to mind.

Ordinarily SAPIENS would accept boy, smash, and bus as one simultaneous, parallel input string. Thus the syntactic features of subject and predicate, or actor and object, are lost. That is, boy, smash, bus and bus, smash, and boy produce exactly the same RFE. Since sentences (3a) and (3b) actually have different meanings, it would be nice if the relevant syntactic information could somehow be preserved.

One way in which this can be done is to give a greater weight to the agent of the input sentence. The agent (subject) might, for example, be assigned an input weight of 10, and the verb and patient might be assigned weights of, say, 1. Such weights could be regarded as reflecting the salience of each case role. Figure 6 shows the effects on the RFE that such a manipulation of weights has. As one might expect, boy-related items are ranked high when the input seed boy

is weighted high, and bus-related items are ranked high when the input seed bus is weighted high.

---

Insert Figure 6 about here

---

This result suggests that the salience of the input words can play an important role in determining the activation levels of nodes in an RFE. The activation levels are important because the nodes with the most activation tend to be more relevant than lower activated nodes.

Needless to say, assigning weights of 10 to the agent, and 1 to the verb and the patient is rather an arbitrary way of handling the problem, yet it seems to work for all that. A more sophisticated approach would be to select the weights on the basis of some kind of optimization procedure. Since SAPIENS has no parsing ability, the user has to apply the weights to each case role. For handling natural language it would be necessary to recognize, for example, that the input was in passive form so as to permit appropriate adjustments to the weights to be made. The purpose of the present demonstration, however, is only to show that, in principle, SAPIENS has the flexibility to be sensitive to simple syntactic constraints.

## Conclusion

Natural language processing, like other areas of AI, has to face the problem of how to reduce the search set of candidate representations if it hopes to utilize appropriate ones to facilitate top down processing. This has proved to be a somewhat stubborn problem in the past, and one which becomes increasingly difficult as the size of the data base increases. SAPIENS appears to be a viable general solution to this problem, because it quickly produces a

small set of candidates and is able to do so in a manner that would permit it to help in a range of important natural language processing tasks. The principles behind the design of SAPIENS are essentially independent of the particular mechanisms that would be employed in a complete natural language processing system, and, in fact, there is no reason why SAPIENS could not be utilized in other AI domains as well. The chief constraint on SAPIENS is that it has to embody a realistic representation of the associative connections between its nodes. In our implementation the onerous task of assembling these data was done independently.

While spreading activation mechanisms have been proposed (and in AI, occasionally used) before in both psychology and AI, such a mechanism has not been shown to be viable or efficient when applied to a data base of any significant size. SAPIENS is an implemented system, rather than an abstract proposal, and as such, the specific details of its design become important, since it is just such details that distinguish approaches that could only work in principle from those that work in fact.

References

Anderson, J. R. Language, memory, and thought. New York: Wiley, 1976.

Anderson, R. C., & Ortony, A. On putting apples into bottles -- A problem of polysemy. Cognitive Psychology, 1975, 7, 167-180.

Bartlett, F. C. Remembering. Cambridge: Cambridge University Press, 1932.

Collins, A. M., & Loftus, E. F. A spreading activation theory of semantic processing. Psychological Review, 1975, 82, 407-428.

Collins, A. M., & Quillian, M. R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247.

Hewitt, C. Description and theoretical analysis (using schemata) of planner: A language for proving theorems and manipulating models in a robot (AI Tech. Report No. 258). Cambridge, Mass.: M.I.T., 1972.

Kiss, G. R. Word associations and networks. Journal of Verbal Learning and Verbal Behavior, 1968, 7, 707-713.

Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Norman, D. A., & Rumelhart, D. E. Explorations in cognition. San Francisco: Freeman, 1975.

Ortony, A. Remembering, understanding, and representation. Cognitive Science, 1978, 2, 53-69.

Piaget, J. The origins of intelligence in children. New York: International Universities Press, 1952.

Quillian, M. R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, Mass.: M.I.T. Press, 1968.

Raphael, B. The thinking computer. San Francisco: Freeman, 1976.

Robinson, J. A. A machine-oriented logic based on the resolution principle. Journal of the Association for Computing Machinery, 1965, 12, 23-41.

Robinson, J. A. The generalized resolution principle. In D. Michie (Ed.), Machine Intelligence 3. Edinburgh: Edinburgh University Press, 1968.

Rosch, E. H. On the internal structure of perceptual and semantic categories. In T. F. Moore (Ed.), Cognitive development and the acquisition of language. New York: Academic Press, 1973.

Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro & W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N.J.: Erlbaum, 1977.

Schank, R. Dynamic Memory. New York: Cambridge University Press, 1982.

Schank, R. C., & Abelson, R. P. Scripts, plans, goals and understanding. Hillsdale, N.J.: Erlbaum, 1977.

Smith, E. E., Shoben, E. J., & Rips. Structure and process in semantic memory: A featural model for semantic decisions. Psychology Review, 1974, 81, 214-241.

Waltz, D. Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Table 1

Results of Disambiguation Simulation

| Context Word | Target Word | First Three Rank-Ordered Nodes in the RFE[a] | Number of Nodes in RFE | Strength Sum | Proportion of Total Activation |
|---|---|---|---|---|---|
| chocolate | mint | sweet, milk, taffee | 3 | 22 | .07 |
| coin | mint | money, cash, gold | 11 | 77 | .26 |
| river | bank | river, water, stream | 9 | 87 | .29 |
| money | bank | money, silver, book | 23 | 123 | .41 |
| metal | bar | iron, door, rod | 3 | 23 | .07 |
| drink | bar | beer, drink, pint | 7 | 83 | .28 |
| yellow | fruit | orange, green, bananna | 4 | 25 | .08 |
| red | fruit | apple, green, hot | 3 | 31 | .10 |
| card | game | play, poker, board | 3 | 28 | .09 |
| ball | game | play, football, tennis | 3 | 33 | .11 |
| chess | game | board, play | 2 | 40 | .13 |

[a]Relevant Forward Environment

Table 2

Results of Cued Recall Simulation

| Target Sentence Pair[a] | Proportion of Total Activation With | |
|---|---|---|
| | Close Cue | Remote Cue[b] |
| | *shampoo* | *detergent* |
| The *nurse* *washed* her *hair*. | .36 | .18 |
| The *nurse* *washed* her *clothes*. | .18 | .22 |
| | *sexy* | *doctor* |
| *Nurses* are often *beautiful* | .23 | .17 |
| *Nurses* give health *care*. | .09 | .51 |
| | *saw* | *scissors* |
| The *farmer* *cut* the *wood*. | .25 | .18 |
| The *farmer* *cut* the *fabric*. | .08 | .29 |
| | *climbing* | *hanging* |
| The *ivy* *covered* the *walls*. | .20 | .16 |
| The *picture* *covered* the *walls*. | .11 | .14 |
| | *hammer* | *fist* |
| The *man* *hit* the *nail*. | .20 | .11 |
| The *man* *hit* the *jaw*. | .07 | .19 |
| | *tv* | *radio* |
| The *man* *watched* the *show*. | .11 | .01 |
| The *man* *listened* to the *show*. | .13 | .11 |

[a] Words in italics represent input seeds.

[b] As presented, the close cue usually produces a higher proportion of total activation for the first member of the pair and a lesser proportion for the second. Similarly, the remote cue produces a lesser proportion for the first member and a higher proportion for the second.

## Figure Captions

Figure 1. Expansion of nodes apple and fruit.

Figure 2. red and pie are intersections found in the expanded environment of apple and fruit.

Figure 3. Original seeds, in this example the nodes apple and fruit, are always weighted at 1.0. Intersections, here the nodes red and pie, are proportionally weighted at fractions which reflect the proportion of total activation on each node.

Figure 4. The end result after one time-slice is a relevant forward environment (RFE).

Figure 5. Steps involved in creating the seed-are intersection used in the cued recall experiment simulation.

Figure 6. Result of simulation of syntactic effects.

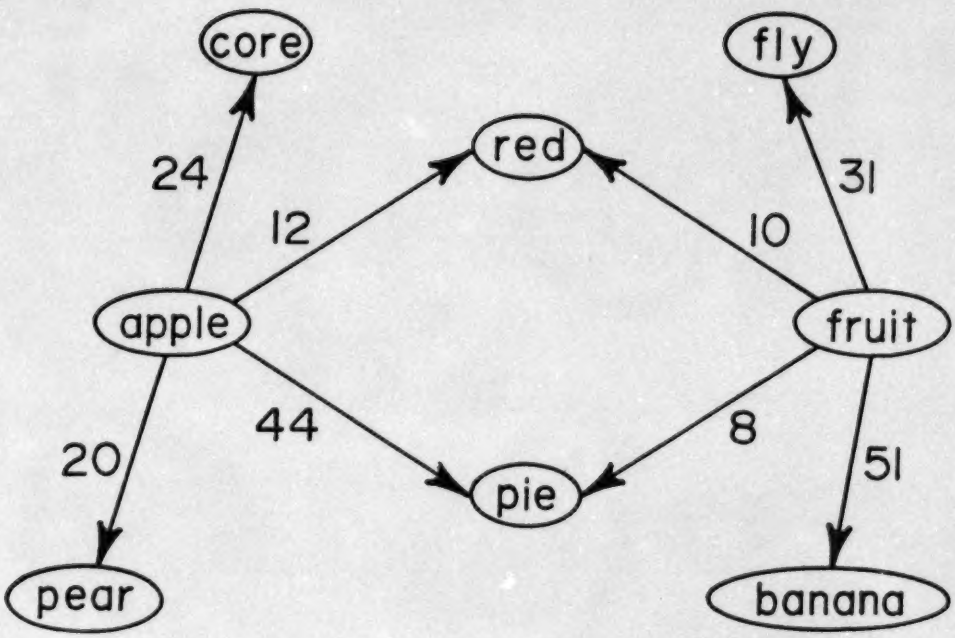
core
fly
red
24
12
10
31
apple
fruit
44
8
20
51
pie
pear
banana

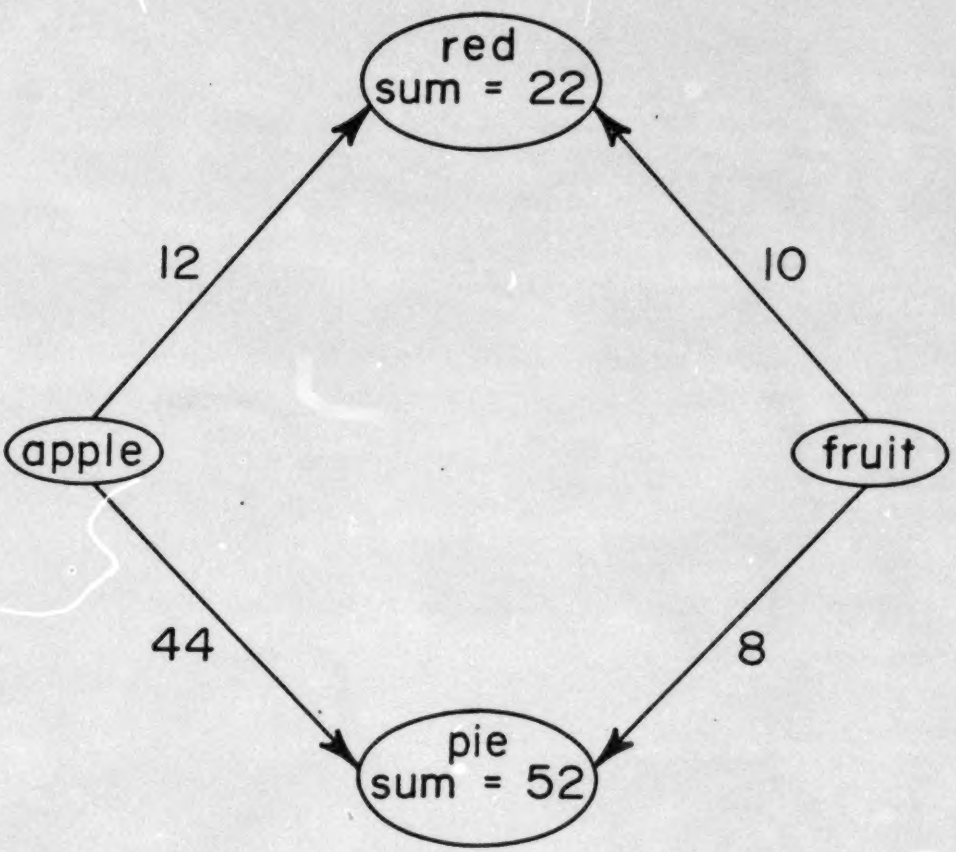
red
sum = 22
12
10
apple
fruit
44
8
pie
sum = 52

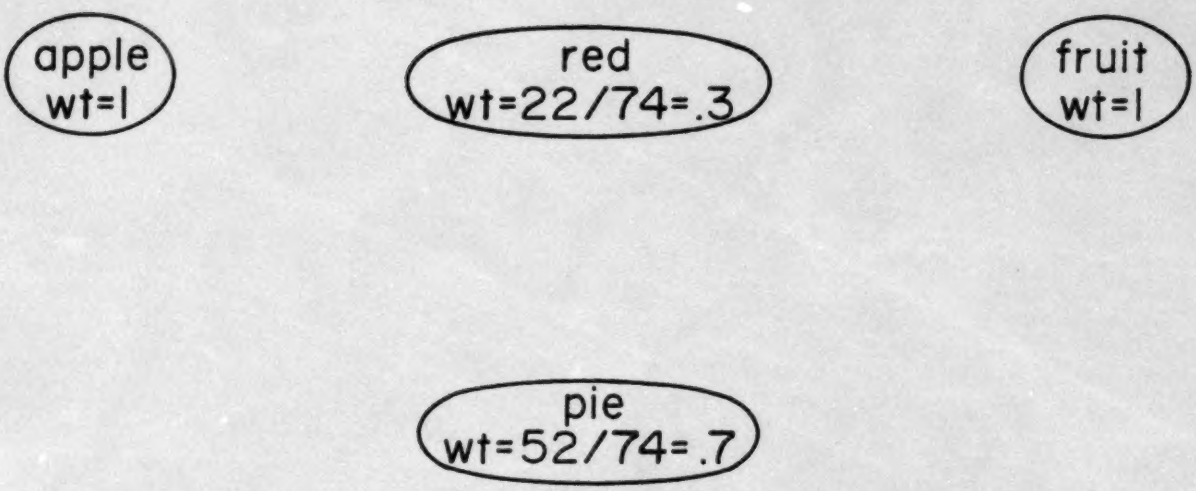
apple
wt=1
red
wt=22/74=.3
fruit
wt=1
pie
wt=52/74=.7

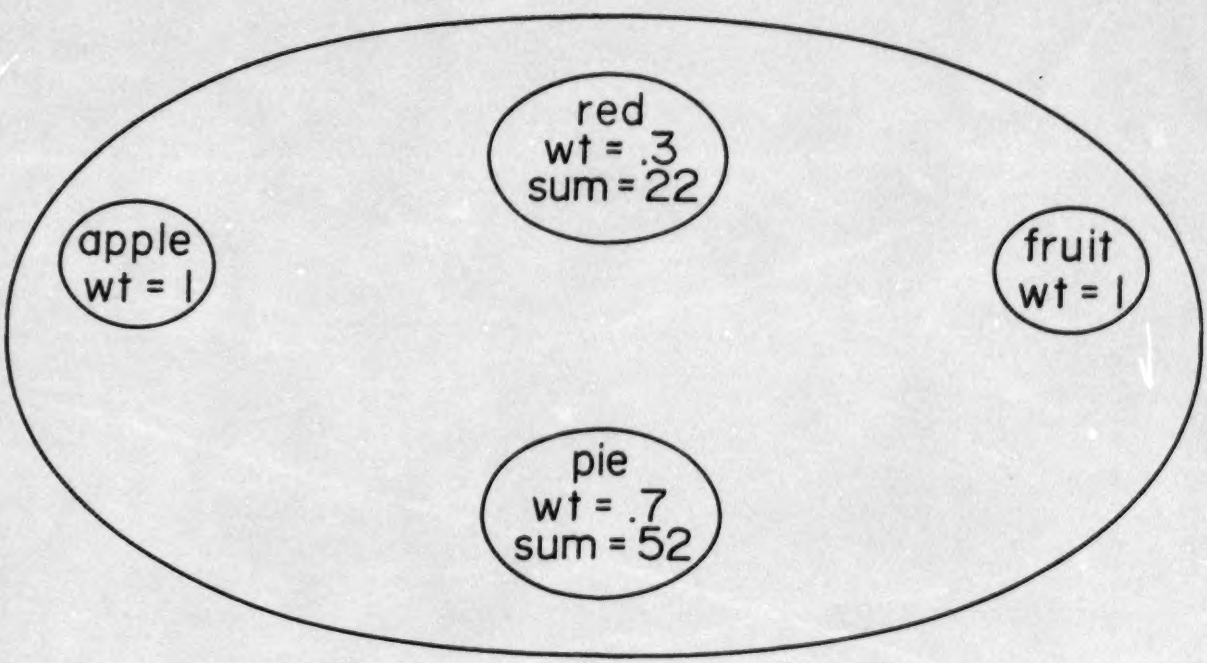
red
wt = .3
sum = 22
apple
wt = 1
fruit
wt = 1
pie
wt = .7
sum = 52

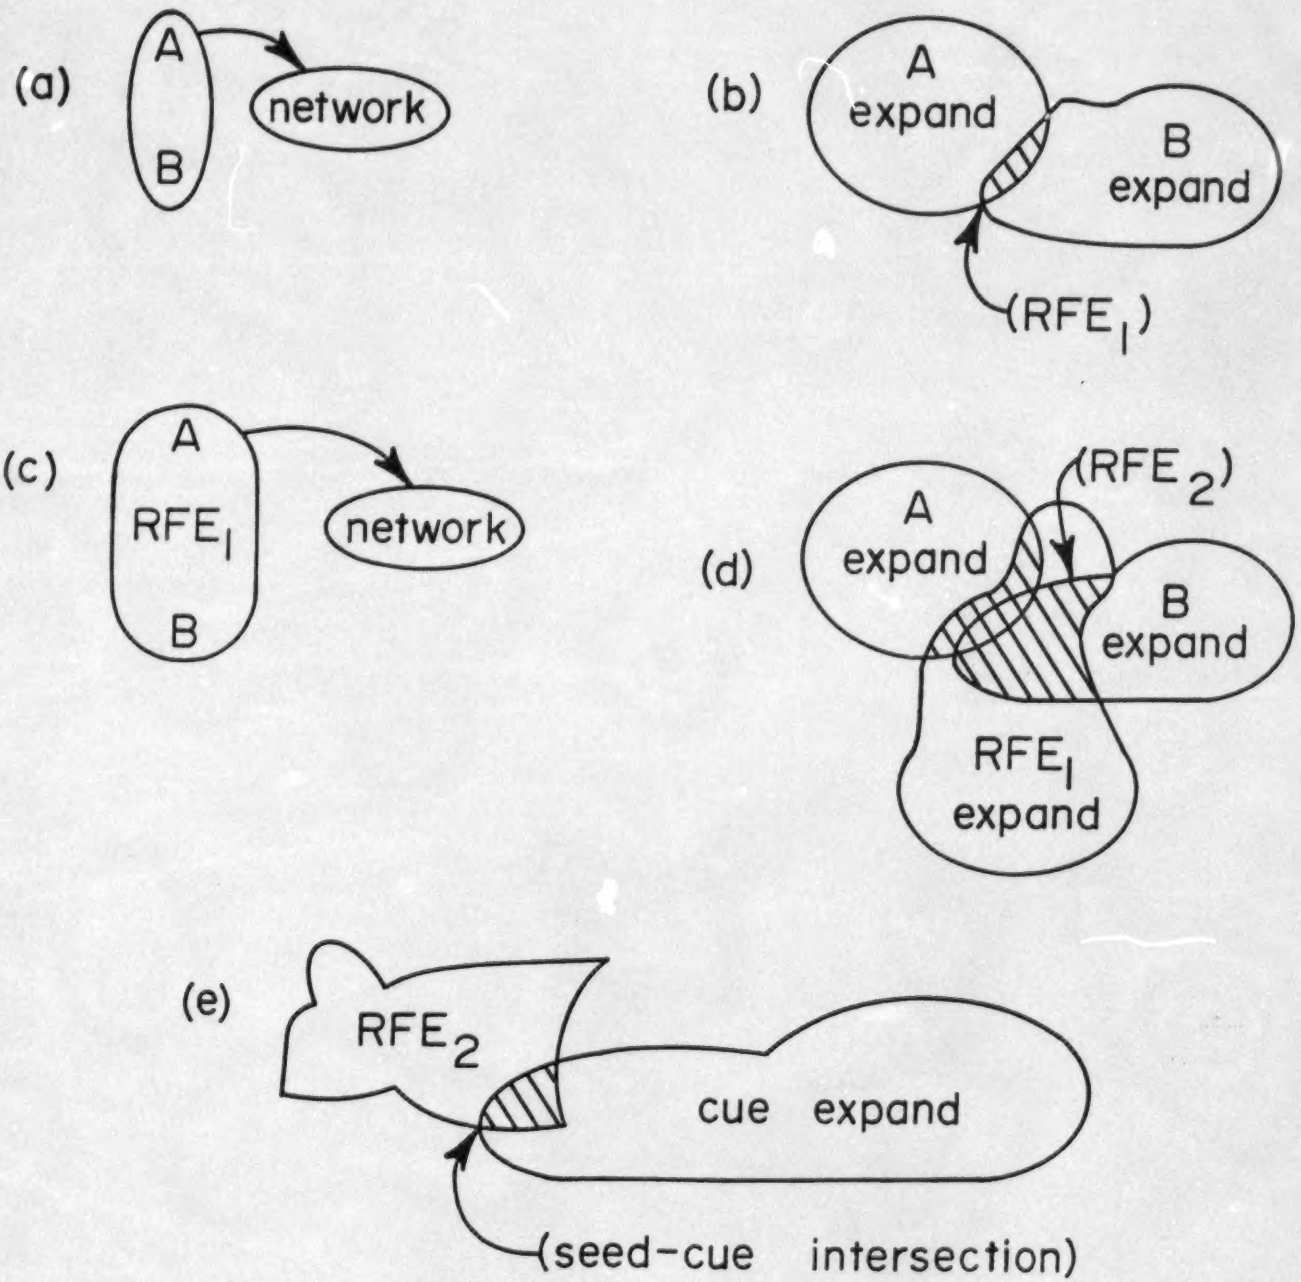
(a)
A
B
network
(b)
A
expand
B
expand
(RFE1)
(c)
A
RFE1
B
network
(d)
(RFE2)
A
expand
B
expand
RFE1
expand
(e)
RFE2
cue expand
(seed-cue intersection)

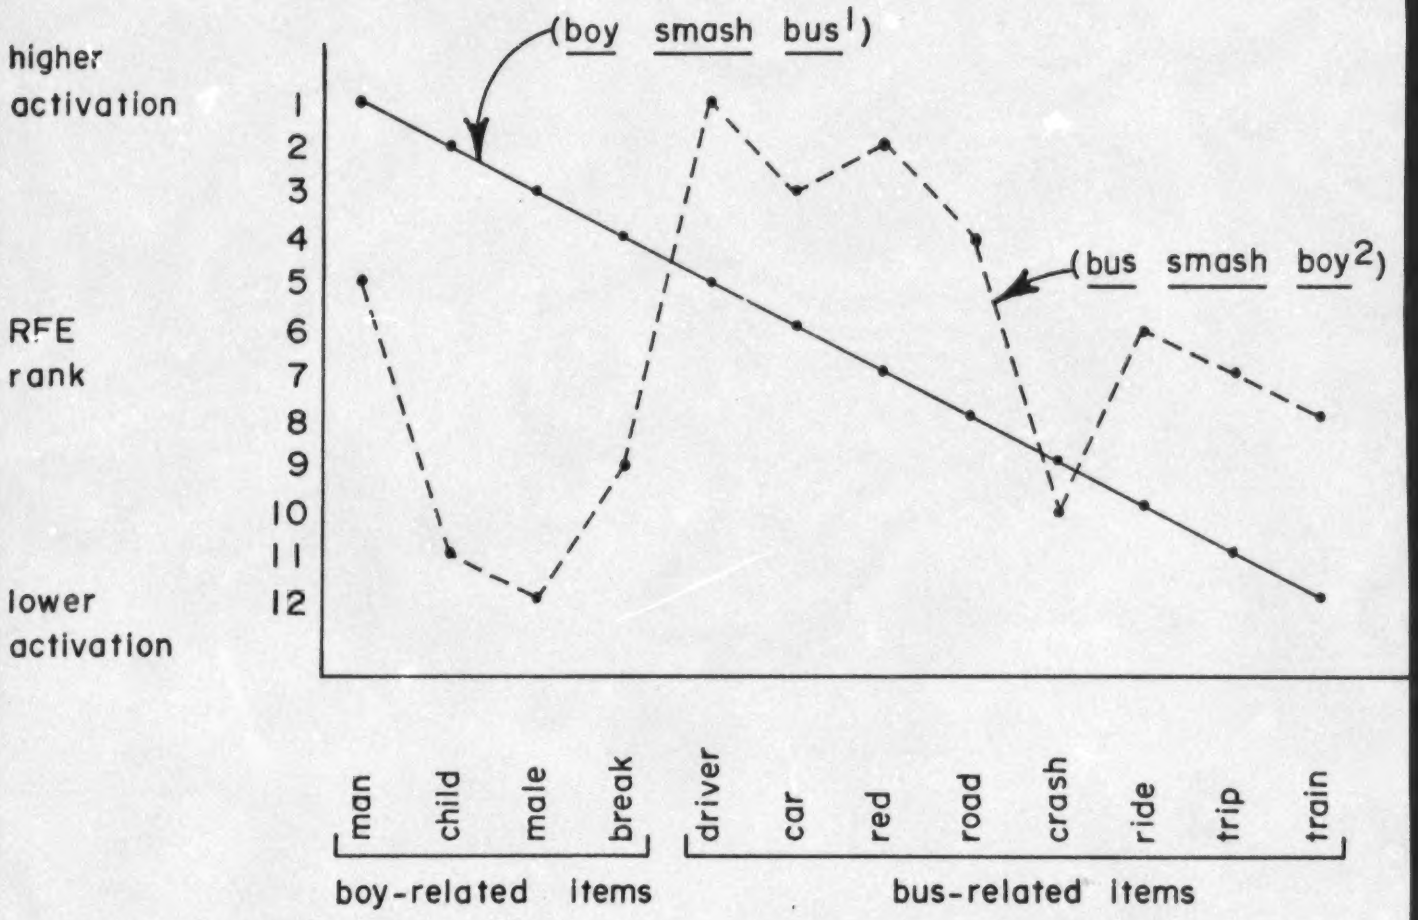


1 This straight line shows the ranking of nodes in the RFE produced by boy smash bus.

2 This curve shows how the nodes are ranked lower for boy-related items and higher for bus-related items when the input seeds are bus smash boy.